CRD report

News from the LBNL Computational Research Division

January 2005

LBNL Speaker Series in Washington to Feature CRD Expertise

Beginning in February, scientists from CRD will launch a series of presentations at Berkeley Lab's project office in Washington, D.C. The goal of the series is to better inform the Washington research community about the achievements and expertise of LNBL staff. The LBNL office is located at 901 D Street, SW, Suite 950. The office is in the Aerospace Center, across D Street from L'Enfant Plaza. Directions to the office can be found at http://www.LBL.gov/Washington/wash-map-page.html.

Wes Bethel

Wes Bethel, leader of the Visualization Group in CRD, will give the first talk at 9 a.m. Thursday, Feb. 17. His presentation is entitled, "Finding the Unknown in a Sea of Data: Leveraging Human Intuition with Scientific Visual Data Analysis."

In his talk, Bethel will describe several approaches to help researchers tackle the problem of detecting patterns in their increasingly complex data and possibly even finding the unexpected. One example is ProteinShop, a software application used for predicting the 3D shape of proteins from amino acid sequences. ProteinShop takes advantage of human intuition to accelerate the protein structure prediction process, thereby reducing the time required from weeks to hours. Another example is a combination of visualization technologies and data management to produce what is

(continued on page 2)

What is CRD Report?

CRD Report is published every other month, highlighting recent achievements by staff members in Berkeley Lab's Computational Research Division. Distributed via email and posted on the Web at http://crd.lbl.gov/DOEresources, CRD Report may be freely distributed. CRD Report is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

FastBit: An Efficient Indexing Technology for Billions of Objects

Three members of the Scientific Data Management Group - John Wu, Arie Shoshani and Ekow Otoo - have been granted a patent for their "Word Aligned Bitmap Compression Method and Data Structure." This technology is currently used in a software called FastBit to compress bitmap indices. When answering a user query, FastBit is often 10 times as fast the searching method used by one of the leading commercial database management systems.

One of the first applications that employs FastBit is in the STAR high-energy physics experiment at Brookhaven National Laboratory. FastBit, which uses the patented Word-Aligned Hybrid (WAH) compression method, has also been applied to spatio-temporal data, such as the direct numerical simulation of a hydrogen-air mixture (Jacqueline Chen, Sandia National Laboratories) and the Terascale Supernova Initiative's simulation of supernova explosions (Anthony Mezzacappa, Oak Ridge National Laboratory).

"Our tests on a set of high-energy physics data show that WAH-compressed bitmap indices are often an order of magnitude faster than the best-known bitmap indexing schemes in commercial systems," according

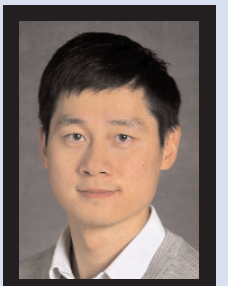

John Wu

to John Wu, the lead developer of FastBit.

Indexing methods such as FastBit and B+tree are used extensively by database management systems to provide fast processing time for user queries. To answer queries on read-only data, such as those from many scientific applications, bitmap indices are known to be more efficient than B+trees. However, without compression, the bitmap indices may require too much space to be useful. While compression reduces the storage requirement, it usually makes query processing slower. To make query processing more efficient, a number of specialized compression schemes have been proposed, with the best-known one called the Byte-aligned Bitmap Code (BBC).

By using a compressed index that can be search without its full decompression, Fastbit uses significantly less space than conventional indexing methods, but is still very efficient.

(continued on page 3)

Face-to-Face Discussion Helps Fusion Scientists Solve Code Interface Problem

Sometimes, $14 can go a long way. For the price of a train ticket from Manhattan to Princeton, CRD's Sherry Li was able to meet with scientists at the Princeton Plasma Physics Lab and together they were able to solve problems that were keeping a new fusion code from running fully parallel.

Li, a member of the Scientific Computing Group and one of the key developers of the SuperLU library of solvers, had been consulting with Steve Jardin's group at PPPL for several months as the fusion researchers worked to develop a newer, faster version of their legacy code known as M3D.

Sherry Li

"We finally decided it might be better to sit down in person and look over the code." - Sherry Li, CRD

M3D used the explicit method to solve partial differential equations, an approach that required many small time steps, which too-longer to run. The new version, called M3D-C1, uses an implicit scheme with much larger time steps, therefore requireing fewer time steps to the solution. "However, the matrix is much more difficult to solve, and many solvers cannot solve it," Li said.

(continued on page 2)

Speaker Series (continued from p.1)

known as "query-driven" visualization. The main idea is to perform intelligent, user-guided selection to find and analyze only the "interesting data," rather than a brute force method to perform analysis on the entire dataset.

Victor Markowitz

The March presentation will feature Victor Markowitz talking about the Biological Data Management and Technology Center (BDMTC). Markowitz' talk will be at 9 a.m. Friday, March 11. Before rejoining the Lab in January 2004, Markowitz was CIO and Senior VP, Data Management Systems at Gene Logic, where he was responsible for the development and deployment of the data management and analysis platform for the company's gene expression data. Prior to joining Gene Logic in 1997, he was a staff scientist at LBNL, where he led the development of data management tools applied to biological databases.

The BDMTC is based on the premise that effectively addressing biological data management challenges requires consolidating data management and system development expertise in a central core. Biological data management involves data generation and acquisition, data modeling, data integration and data analysis. Data management challenges are posed by increasing amounts of experimental data generated by life science applications, difficulty of qualifying data generated using inherently imprecise tools and techniques, and complexity of integrating data residing in diverse and poorly correlated repositories. Biological data management systems in academic settings such as LBNL are developed with minimal or no user and system analysis, without following system development practices, and without considering system evolution, maintenance and scalability.

This presentation will discuss the main challenges of biological data management, the problems encountered by academic groups in addressing these challenges, the rationale for BDMTC, its activities to date and future plans.

For more information about the series, contact Jon Bashor at JBashor@lbl.gov.

News About the ACTS Collection

In conjunction with demonstrations at the SC2004 conference last November, Osni Marques and Tony Drummond of CRD's Scientific Computing Group created a matrix of applications that use ACTS tools. Marques and Drummond provide support for the DOE ACTS (Advanced CompuTational Software) Collection, a set of software tools developed by DOE, sometimes in collaboration with other agencies (DARPA, NSF), that make it easier for programmers to write high performance scientific applications for parallel computers.

The purpose of the matrix is to present a summary of the utilization of ACTS tools in important scientific and engineering applications from an international pool of users. The goal is to showcase not only the benefits from the use of robust and widely used software tools, but also to provide guidance on how software tools work in different application areas. Please see http://acts.nersc.gov/MatApps for more information.

The ACTS project team has also prepared a CD containing information about the tools currently available in ACTS. The CD contains copies of the tutorials and talks presented at the Fifth ACTS Collection Workshop held at LBNL on August 24-27, 2004, the matrix of applications that use ACTS tools referred to above, and more. To receive a copy of this CD and learn more about what is going on with ACTS, send an e-mail to acts-support@nersc.gov with your name and mailing address.

Discussion Helps Fusion Scientists *(continued from page 1)*

Several months ago, the team began using SuperLU as their solver. Although the fusion group computes on Seaborg, the IBM supercomputer at NERSC, they had acquired a 16-processor SGI Altix system to do local development of their new code.

As problems arose, Li and the PPPL team exchanged emails trying to resolve the problems. They even sent their code to her, but she could not find the sticking points just by reviewing it.

"It was getting more complicated – they know the physics part and I know the solver part," Li said. "We finally decided it might be better to sit down in person and look over the code."

So, while attending the 16th International Conference on Domain Decomposition Methods at NYU's Courant Institute in January, Li slipped away for a day and took the train to Princeton.

"They educated me more on how their code worked and we looked at the interface," she said. "M3D is written in Fortran 90 and my code is in C, so we needed to build some new wrappers." In the process of debugging in real time, they were able to identify a word-type inconsistency in the interface that caused the SGI implementation to fail for the largest problem sizes.

"Even though you have been very responsive via email during the last few months, there was really no substitute for your actually being here to witness and diagnose the problems we were having," Jardin wrote to Li after their meeting. "Thank you so much for making a special trip from your conference to help us debug the implementation of your distributed SuperLU software on our local SGI Altix. This has really made a big impact. As a result of your visit, not only do we understand your SuperLU-dist much better, but we are now able to run our largest jobs in a fully parallel mode, with even better than 'ideal' scaling. This will really help us in our code-development activities for the new M3D-C1 code, and will also make our use of NERSC for this code much more productive."

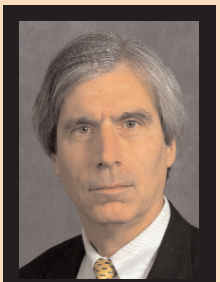

Steve Jardin

"As a result of your visit, not only do we understand your SuperLU-dist much better, but we are now able to run our largest jobs in a fully parallel mode, with even better than 'ideal' scaling." - Steve Jardin, PPPL

While the immediate results demonstrate the value of an interpersonal collaborative approach, the success is also an example of how DOE's SciDAC program is meeting its goal of developing advanced tools though collaboration. The SuperLU development is partly funded by the TOPS SciDAC project led by David Keyes of Columbia University, while M3D-C1 is funded by the fusion CEMM SciDAC project, which is led by Jardin.

FastBit Speeds up Searches *(continued from page 1)*

2.2 million records 12 common attributes		B+tree DBMS	Projection index	Bitmap index		
				DBMS	BBC	WAH
Index size (MB)		408	113	111	118	186
Average query processing time (seconds)	1 attribute	0.95	0.51	0.05	0.03	0.02
	2 attributes	2.15	0.56	0.39	0.17	0.04
	5 attributes	2.23	0.67	2.42	0.76	0.17

As can be seen from the table above, FastBit's WAH indexing method creates an index that is slightly larger than other compressed bitmap indexing methods, but the time needed to process a query is less, often much less. "It's really a space-time tradeoff," Wu said.

FastBit is currently being used by several DOE research projects and has yielded several success stories:

- GridCollector, the software module for the STAR analysis framework, uses FastBit's searching capability to provide STAR analysts with a new way of accessing collision data called “events.” Instead of naming the data files as was previously done, analysts can now select events based on physically meaningful attributes known as tags. Through GridCollector, analysis programs only read the selected events, instead of every event in the selected data files. Since most analysis jobs use only a small fraction of the events in data files, the GridCollector can significantly improve the turnaround time.

- Tracking features in the analysis of combustion simulation data is more efficient. By using FastBit and compressed bitmaps, the FastBit team was able to significantly speed up the problem of tracking ignition kernels in a combustion simulation of a hydrogen-air mixture. This approach addressed the difficult problem of identifying data indicating the ignition kernel from the rest of the simulation data and tracking the progression of flames over time.

- DEX, or Dexterous Data Explorer, is a collaboration between the Scientific Data Management Group and the Visualization Group in CRD. DEX uses FastBit to provide query-based visualization of large scientific datasets. A preliminary version of the software was demonstrated at the Supercomputing 2004 conference on both combustion datasets and supernova simulation datasets. LBNL collaborators on DEX are Kurt Stockinger, John Shalf and Wes Bethel.

- Compressed bitmaps are also used in a view-dependent isosurface software. At Supercomputing 2004, a preliminary version of the software was demonstrated to display in real-time the isosurfaces of large complex data produced from a simulation of Richtmyer-Meshkov Instability, which is the impulsive-acceleration limit of the Rayleigh-Taylor instability in computational fluid dynamics. In this application, compressed bitmaps are used to record what data items are visible from a particular viewing angle. This allows the software to extract the minimal amount of data items required for visualization and make it very efficient to render very large complex isosurfaces as the user changes viewing angles.

The effectiveness of FastBit has also attracted the attention of other institutions as well. The ROOT developers at CERN have started working on incorporating FastBit into their software. Since ROOT software is used by most major high-energy physics projects around the world, fully integrating FastBit into it would far extend the user community in the scientific realm. There is also interest from the commercial sector.

“We have even learned that at least one foreign telecommunications providers has implemented WAH compression for their data analysis,” Wu said.

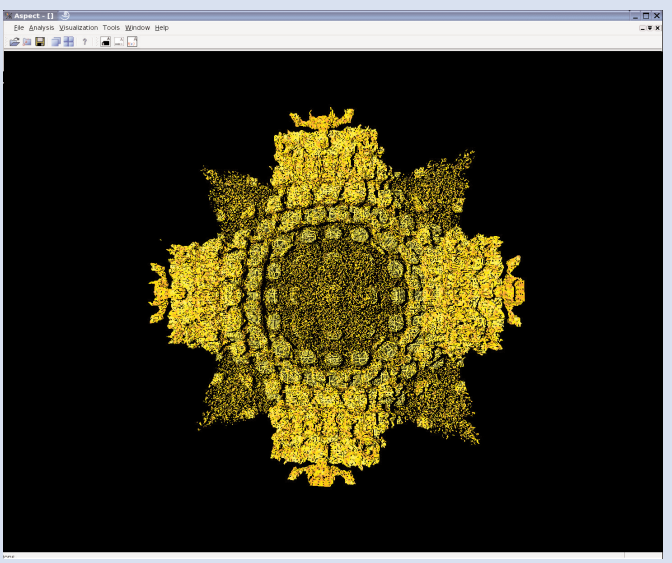

An isosurface produced by the view-dependent isosurface software. As shown, the isosurface software is run under an analysis framework called ASPECT developed at Oak Ridge National Laboratory.

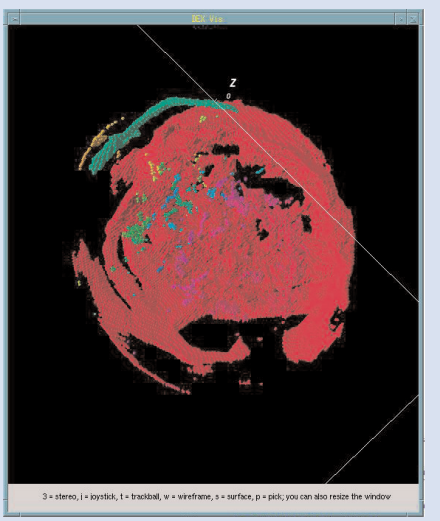

Regions of interest in a Supernova simulation dataset, selected with the DEX software.